

Cisco Overlay Transport Virtualization

- Q.** What is Cisco Overlay Transport Virtualization?
- A.** Cisco® Overlay Transport Virtualization (OTV) is a “MAC in IP” technique for supporting Layer 2 Vans over any transport.
- Q.** What are the OTV transport requirements for the core?
- A.** The overlay nature of OTV allows it to work over any transport as long as this transport can forward IP packets. Any optimizations performed for IP in the transport will benefit the OTV encapsulated traffic.
- Q.** Does OTV require the configuration of pseudo-wires or tunnels?
- A.** OTV does not require the configuration of pseudo-wires or tunnels. The encapsulation of the packets sent to the overlay is performed dynamically based on a Layer 2 address destination lookup.
- Q.** Is OTV a public standard?
- A.** OTV is not a public standard, even though the technology is based on standardized protocols. The plan is to submit the IETF drafts toward the end of calendar year 2010 (CY2010).
- Q.** On which Cisco platforms will OTV be supported?
- A.** OTV will make its first appearance on the Cisco Nexus® 7000 Series Switches. Other platforms in the Cisco portfolio are planning to support this technology as well.
- Q.** What modules of the Cisco Nexus 7000 Series will support OTV?
- A.** OTV is supported on all M-series modules. OTV is not supported on F1-series modules. Deployments using F1 series can leverage VDC separation to achieve the desired combination of line cards and functionality.
- Q.** Does OTV affect the existing Layer 2 design at a site?
- A.** OTV is both core and site transparent. No changes to the Layer 2 design of the sites are needed.
- Q.** Does OTV require the extension of the Spanning Tree Protocol?
- A.** OTV can extend the Layer 2 domains across geographically distant data centers by providing built-in filtering capabilities to localize the most common networking protocols (Spanning Tree Protocol, VLAN Trunking Protocol [VTP], and Hot Standby Router Protocol [HSRP]) and prevent them from traversing the overlay, therefore keeping protocol failures from propagating across sites.
- With OTV, each site can have its own implementation of spanning tree as well as its own spanning tree root device.
- Q.** Does OTV provide site localization for the First-Hop Routing Protocol (FHRP)?
- A.** OTV provides built-in filtering capabilities to localize FHRP. A single FHRP group across the extended Layer 2 domain will now have an active gateway on each site, providing optimal egress routing.
- Q.** How does OTV propagate the MAC addresses learned at each site?
- A.** Unlike traditional Layer 2 Vans, which rely on Layer 2 flooding to propagate MAC address reachability, OTV uses a protocol to proactively advertise the MAC addresses learned at each site. The protocol advertisement takes place in the background, with no configuration required by the network administrator.

-
- Q.** How does OTV handle unknown Unicast flooding?
- A.** OTV does not require unknown Unicast flooding to propagate MAC address reach ability; thus, there is no need to flood the unknown Unicast over the overlay, and flooding is suppressed. The endpoints connected to the network are assumed to be neither silent nor unidirectional. OTV also provides a way to learn the MAC addresses for unidirectional hosts.
- Q.** Does OTV support static MAC-to-IP address mapping?
- A.** OTV enables the network administrator to statically map MAC addresses to IP addresses.
- Q.** What are the OTV advantages for multicast traffic?
- A.** With OTV, the multicast traffic generated at the site is optimally replicated by the core. Head-end replication is not performed at the site when the core provides multicast services. OTV encapsulates the multicast frame into a multicast IP packet (after joining a multicast group in the core), which will be replicated by the core to only those sites that are supposed to receive it. This capability is built-in to OTV and does not require any multicast configuration at the sites.
- Q.** Does OTV require multicast support in the core?
- A.** Although desirable to fully benefit from all the optimizations that OTV can provide, multicast in the core is not mandatory. *Starting with NX-OS 5.2* OTV also provides an elegant and dynamic solution for those cores that do not have multicast support.
- Q.** Does OTV support multiple overlays?
- A.** OTV can support multiple concurrent overlays.
- Q.** Does OTV support multipoint?
- A.** Automatic detection of multipoint is included as part of the OTV control protocol. This feature enables multipoint of sites without requiring any additional configuration or protocol.
- Q.** How does OTV provide load balancing?
- A.** OTV provides two levels of load balancing:
- Within the core: OTV headers are defined to allow the core to hash traffic based on five-tuples and distribute traffic over multiple paths to avoid polarization of encapsulated traffic.
 - Within the site: OTV enables effective load balancing of flows across the multiple edge devices available in an all-active multihomed deployment. Load balancing follows equal-cost multipath (ECMP) rules based on the information provided by the OTV control protocol.
- Q.** What is the encapsulation used by OTV?
- A.** OTV uses Ethernet over Generic Router Encapsulation (GRE) and adds an OTV shim to the header to encode VLAN information. The OTV encapsulation is 42 bytes, which is less than virtual private LAN service (VPLS) over GRE. The encapsulation is performed entirely by the forwarding engine in hardware.
- Q.** How complicated is OTV configuration?
- A.** OTV has been designed with only a few command-line interface (CLI) and built-in automated processes. This design makes OTV extremely simple to configure.



Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV Amsterdam,
The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at www.cisco.com/go/offices.

Cisco and the Cisco Logo are trademarks of Cisco Systems, Inc. and/or its affiliates in the U.S. and other countries. A listing of Cisco's trademarks can be found at www.cisco.com/go/trademarks. Third party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1005R)

Printed in USA

C67-574969-02 09/12